

Claims:

1. A method for processing a speech signal comprising:
extracting prosodic features from a speech signal;
modeling the prosodic features to identify at least one speech endpoint; and
producing an endpoint signal corresponding to the occurrence of the at least one speech endpoint.
2. The method of claim 1 wherein the extracting step comprises:
processing pitch information within the speech signal.
3. The method of claim 2 wherein the extracting step further comprises:
determining a duration pattern; and
performing pause analysis.
4. The method of claim 2 wherein the processing step comprises:
generating a pitch contour;
producing a pitch movement model from the pitch contour; and
extracting at least one pitch parameter from the pitch movement model.
5. The method of claim 4 wherein the at least one pitch parameter is a pitch movement slope.
6. The method of claim 4 wherein the at least one pitch parameter is a difference between the pitch information in the speech signal and baseline pitch information.
7. The method of claim 1 wherein the producing step comprises generating a posterior probability regarding the at least one speech endpoint.
8. The method of claim 7 wherein the posterior probability regarding a plurality of speaker states including a probability that a speaker has completed an utterance, a probability that the speaker is pausing due to hesitation, or a probability that the speaker is talking fluently.

9. The method of claim 8 where the posterior probability is continuously updated as the speech signal is processed.
10. The method of claim 1 further comprising:
 executing a speech recognition routine for processing the speech signal using the at least one speech endpoint.
11. Apparatus for processing a speech signal comprising:
 a prosodic feature extractor for extracting prosodic features from the speech signal;
 a prosodic feature analyzer for modeling the prosodic features to identify at least one speech endpoint; and
 an endpoint signal producer that produces an endpoint signal corresponding to the occurrence of the at least one speech endpoint.
12. The apparatus of claim 11 wherein the prosodic feature extractor comprises:
 a pitch processor for processing pitch information within the speech signal.
13. The apparatus of claim 12 wherein the prosodic feature extractor further comprises:
 means for determining a duration pattern; and
 means for performing pause analysis
14. The apparatus of claim 12 wherein the pitch processor comprises:
 means for generating a pitch contour;
 means for producing a pitch movement model from the pitch contour; and
 means for extracting at least one pitch parameter from the pitch movement model.
15. The apparatus of claim 14 wherein the at least one pitch parameter is a pitch movement slope.
16. The apparatus of claim 14 wherein the at least one pitch parameter is a

difference between the pitch information in the speech signal and baseline pitch information.

17. The apparatus of claim 11 wherein the endpoint signal producer comprises a posterior probability generator for generating a posterior probability regarding the at least one speech endpoint.

18. The apparatus of claim 17 wherein the posterior probability regarding a plurality of speaker states includes a probability that a speaker has completed an utterance, a probability that the speaker is pausing due to hesitation, or a probability that the speaker is talking fluently.

19. The method of claim 18 where the posterior probability is continuously updated as the speech signal is processed.

20. The method of claim 11 further comprising:

a computer for executing a speech recognition routine for processing the speech signal using the at least one speech endpoint.

21. An electronic storage medium for storing a program that, when executed by a processor, causes a system to perform a method for processing a speech signal comprising:

extracting prosodic features from a speech signal;
modeling the prosodic features to identify at least one speech endpoint; and
producing an endpoint signal corresponding to the occurrence of the at least one speech endpoint.